

A SYSTEM FOR MANAGING NETWORKS USING RULES AND INCLUDING AN
INFERENCE ENGINE
CROSS-REFERENCE TO RELATED APPLICATIONS

This application is based on French Patent Application No. 02 09 741 filed July 31, 2002, the disclosure of which is hereby incorporated by reference thereto in its entirety, and the priority of which is hereby claimed under 35 U.S.C. §119.

BACKGROUND OF THE INVENTION

Field of the invention

The present invention relates to the field of telecommunications services management. To be more precise, it relates to the management of telecommunication services using policy rules. The invention applies particularly well to networks using protocols of the Internet Protocol (IP) family or other protocols of higher level.

Description of the prior art

Such networks provide services of various types, including virtual private networks (VPN), videoconferences, etc.

The provision of these services impacts on the behavior expected of the network. The expected behavior can include compliance with a particular quality of service (QoS) associated with the service. In this case, the quality of service is negotiated between at least the operator of the telecommunication network and the provider of the service, in the form of a service level agreement (SLA). The SLA is then specified in a more technical form in a service level specification (SLS), which can conform to the specifications of the Internet engineering task force (IETF).

In other words, the SLS is derived from an SLA and contains the technical parameters that must be used to implement the service.

To provide a service on a telecommunication network, it is therefore necessary to set the network parameters to enable the service to be established, including compliance with the negotiated quality of service, for example.

The parameters can be set using policy rules, referred to for simplicity hereinafter as rules. The rules typically include a set of conditions and a set of actions. The sets can be reduced to a single element, i.e. a rule may consist of only one condition and/or only one action.

Figure 1 shows how rules are implemented.

Conventionally, they are defined at the level of a policy manager (PM) and then transmitted to a policy server (PS). The policy server is responsible for their application by network elements which in this context are referred to as policy enforcement points (PEP).

The policy manager and the policy server are conventionally part of the network management layer (NML); to be more precise, they can belong to a network management system (NMS). However, it is important to note that a network may include only one of these elements, as the policy manager PM and the policy server PS can be two independent physical systems that can be marketed separately.

It is apparent that there is an important semantic difference between the definition of the service, for example in the context of an SLA/SLS, and the corresponding rules, which must be implemented by the network elements or PEP, in particular the configurations of the network elements.

In concrete terms, the difference can become apparent at two or more levels:

Firstly, it obliges the designer of the service to have network expert knowledge. For example, it is incumbent on the service designer to decide how a virtual private network VPN should be implemented, for example whether the IPsec protocol must be used, or if preference must be given to the multi-protocol label switching (MPLS) technology.

Secondly, it obliges the service designer to have access to the exact specifications of each network element to be configured. Depending on the manufacturer, the same type of network element (IP router, firewall, etc.) may be configured differently, because the capacities may be different.

The object of the present invention is to alleviate this drawback and to facilitate the development of new services by means of rules.

SUMMARY OF THE INVENTION

To this end, the invention provides a network management system for implementing a service on a network, the system including means for acquiring policy rules for configuring the service, means for determining commands corresponding to the policy rules and transmitting them to network elements, and processing means for inferring the rules in order to

determine the commands, in which system the rules comprise services rules and implementation rules.

In one embodiment of the invention the processing means include an inference engine.

5 In one embodiment of the invention the implementation rules include technology rules and/or equipment rules.

Thus new services can be designed independently of the implementation by adding processing means to the network management system able to infer services rules and implementation rules dynamically.

10 In particular, the design process does not have to take account of the specifics of the various network elements or of expert data to choose between a set of technical solutions for implementing the new services.

The invention and its advantages will become more clearly apparent in the course of the following description of one embodiment of the invention, which refers to the accompanying drawing.

BRIEF DESCRIPTION OF THE DRAWING

Figure 1, already commented on, represents a prior art system for managing a network using policy rules.

20 Figure 2 is a diagram of a network management system according to the invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

Figure 2 shows a network management system NMS associated with a network N. The network N includes network elements $E_1, E_2 \dots E_n$ which can be IP routers, asynchronous transfer mode (ATM) switches, etc.

25 The network management system NMS includes processing means IE and two databases D_T and D_E . Of course, the two databases could be two views of the same physical database.

The processing means IE preferably comprise an inference engine.

30 The processing means receive as input service rules R_S and implementation rules. In the figure 2 example, the implementation rules are technology rules R_T and equipment rules R_E .

A service rule can consist in creating a virtual private network (VPN) during a specified time period, for example.

35 Simplifying, a rule of this kind could take the form: "IF (timeperiod=march 2002) THEN (create VPN from site A to site B)". This rule

stipulates that a virtual private network must be created between sites A and B during March 2002.

5 The processing means IE further employ implementation rules. The implementation rules can contain technology rules R_T , for example, stored in a database D_T .

The technology rules are used to model expert know-how and automate its application.

10 Accordingly, in the above example concerning the provision of a virtual private network, a choice may be made between different technologies. In particular, it can be implemented using the IPsec protocol, as defined in RFC 2401 of the Internet Engineering Task Force (IETF), or using multi-protocol label switching (MPLS) tunnels, as defined in IETF RFC 3031.

15 One strategy for choosing the technology might be to consider the number of sites involved in the virtual private network and to use that number as a basis for deciding which is the most appropriate technology: for example, if the number of sites is less than five, then the IPsec protocol is preferred, whereas otherwise the MPLS protocol is chosen.

20 This strategy can be modeled in the form of technology rules R_T and stored in the technology database D_T .

Simplifying, the technology rules R_T can be written in the following form:

IF (number_of_sites < 5) THEN (tunneling technology = IPsec)

IF (number_of_sites \geq 5) THEN (tunneling technology = MPLS).

25 The processing means IE can then correlate the service rules with the technology rules. The processing means can in particular include an inference engine. Inference engines include the "IlogRules" product from the company Ilog and the Java Expert System Shell (Jess).

30 In the same way, the processing means can use equipment rules R_E , which can be stored in an equipment database D_E .

35 The equipment rules are used to model how the rules must be adapted or selected for a particular equipment type. This is because two network equipments can have different capacities, even if they are functionally identical. Their capacities may depend on the network equipment manufacturer, or differ between different models in the range of

the same manufacturer. For example, some equipment (such as routers) can optionally support the MPLS technology. The equipment rules R_E can take this into account, so that the management system chooses the right implementation.

5 Returning to the same example, an equipment rule R_E can be written as follows:

 IF (equipment = TYPE1) THEN (tunneling technology=IPsec)

 This means that if the Type1 equipments cannot support the MPLS technology, then IPSec is the only option.

10 If the equipment is not of Type1, then in this example no equipment rule is specified and the choice of the right implementation is effected on the basis of the technology rules R_T previously referred to.

 Accordingly, the services can be described in the form of service rules R_S independently of the technology to be used and the specifics of the
15 network equipment. The aspects related to the technology to be used and to those specifics can be modeled in the form of implementation rules (or metarules).